Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2

Effective date of registration: 11/13/2014 | **spanning-tree bridge assurance**

To enable Bridge Assurance on the device, use the spanning-tree bridge assurance command. To disable Bridge Assurance, use the no form of this command.

spanning-tree bridge assurance

no spanning-tree bridge assurance

Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 115. | **spanning-tree bridge assurance**

The spanning-tree bridge assurance command enables bridge assurance on all ports with a port type of *network*. Bridge assurance protects against unidirectional link failure, other software failure, and devices that quit running a spanning tree algorithm.

Bridge assurance is available only on spanning tree *network* ports on point-to-point links. Both ends of the link must have bridge assurance enabled. If the device on one side of the link has bridge assurance enabled and the device on the other side either does not support bridge assurance or does not have it enabled, the bridge assurance enabled port is blocked.

The no spanning-tree bridge assurance command disables bridge assurance.

The spanning-tree bridge assurance and default spanning-tree bridge assurance commands restore the default behavior by removing the no spanning-tree bridge assurance command from *running-config*. Only the no form of this command is visible in *running-config*.

Platform   all
Command Mode   Global Configuration

Command Syntax

spanning-tree bridge assurance
no spanning-tree bridge assurance
default spanning-tree bridge assurance

Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 967. |

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	spanning-tree guard
	
	To enable or disable Loop Guard or Root Guard, use the spanning-tree guard command. To return to the default settings, use the no form of this command.
	
	spanning-tree guard {loop | root | none}
	
	no spanning-tree guard
	
	Syntax Description
	loop — Enables Loop Guard on the interface.
	root — Enables Root Guard on the interface.
	none — Sets the guard mode to none.
	
	Defaults — Disabled
	
	Command Modes — Interface configuration
	
	SupportedUserRoles — network-admin, vdc-admin
	
	Command History
	Release: 4.0 — Modification: This command was introduced.
	
	Usage Guidelines — You cannot enable Loop Guard if Root Guard is enabled, although the device accepts the command to enable Loop Guard on spanning tree edge ports.
	This command does not require a license.
	
	Examples — This example shows how to enable Root Guard:
	`switch(config-if)# spanning-tree guard root`
	`switch(config-if)#`
	
	Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 119.
		spanning-tree guard
		
		The spanning-tree guard command enables root guard or loop guard on the configuration mode interface. The spanning-tree loopguard default command configures the global loop guard setting.
		
		• Root guard prevents a port from becoming a root or blocked port. A root guard port that receives a superior BPDU transitions to the root-inconsistent (blocked) state.
		
		• Loop guard protects against loops resulting from unidirectional link failures on point-to-point links by preventing non-designated ports from becoming designated ports. When loop guard is enabled, a root or blocked port transitions to loop-inconsistent (blocked) state if it stops receiving BPDUs from its designated port. The port returns to its prior state when it receives a BPDU.
		
		The no spanning-tree guard and default spanning-tree guard commands sets the configuration mode interface to the global loop guard mode by removing the spanning-tree guard statement from *running-config*. The spanning-tree guard none command disables loop guard and root guard on the interface, overriding the global setting.
		
		Platform — all
		Command Mode — Interface-Ethernet Configuration / Interface-Port-Channel Configuration
		
		Command Syntax
		`spanning-tree guard PORT_MODE`
		`no spanning-tree guard`
		`default spanning-tree guard`
		
		Parameters
		• *PORT_MODE* the port mode. Options include:
		— loop — enables loop guard on the interface.
		— root — enables root guard on the interface.
		— none — disables root guard and loop guard.
		
		Examples
		• This command enables root guard on Ethernet 5 interface.
		`switch(config)#interface ethernet 5`
		`switch(config-if-Et5)#spanning-tree guard root`
		`switch(config-if-Et5)#`
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1005.

Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	To enable Loop Guard as a default on all ports of a given bridge, use the spanning-tree loopguard default command. To disable Loop Guard, use the no form of this command. Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 121.	• spanning-tree loopguard default command enables loop guard as a default on all switch ports. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 996.

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**spanning-tree loopguard default** To enable Loop Guard as a default on all ports of a given bridge, use the spanning-tree loopguard default command. To disable Loop Guard, use the no form of this command. spanning-tree loopguard default no spanning-tree loopguard default Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 121.	**spanning-tree loopguard default** The spanning-tree loopguard default command configures the global loop guard setting as *enabled*. Ports not covered by a spanning-tree guard command use the global loop guard setting. Loop guard prevents blocked or root ports from becoming a designated port due to failures resulting in a unidirectional link. The spanning-tree guard interface configuration statement overrides the global setting for a specified interface. The default global loop guard setting is *disabled*. The no spanning-tree loopguard default and default spanning-tree loopguard default commands restore the global loop guard setting of *disabled* by removing the spanning-tree loopguard default command from *running-config*. Platform all Command Mode Global Configuration Command Syntax spanning-tree loopguard default no spanning-tree loopguard default default spanning-tree loopguard default Examples • This command enables loop guard as the default on all switch ports. switch(config)#spanning-tree loopguard default switch(config)# Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1008.

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**spanning-tree mst configuration** To enter the Multiple Spanning Tree (MST) configuration submode, use the spanning-tree mst configuration command. To return to the default settings, use the no form of this command. spanning-tree mst configuration no spanning-tree mst configuration Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 124.	**spanning-tree mst configuration** The spanning-tree mst configuration command places the switch in MST-configuration mode, which is the group change mode where MST region parameters are configured. Changes made in a group change mode are saved by leaving the mode through the exit command or by entering another configuration mode. To discard changes from the current edit session, leave the mode with the abort command. These commands are available in MST-configuration mode: • abort (mst-configuration mode) • exit (mst-configuration mode) • instance • name (mst-configuration mode) • revision (mst-configuration mode) • show (mst-configuration mode) The no spanning-tree mst configuration and default spanning-tree mst configuration commands restore the MST default configuration. Platform all Command Mode Global Configuration Command Syntax spanning-tree mst configuration no spanning-tree mst configuration default spanning-tree mst configuration Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1012.
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	Related Commands | Command	Description
---	---
instance vlan	Maps a VLAN or a set of VLANs to an MST instance.
name (mst configuration)	Sets the name of an MST region.
revision	Sets the revision number for the MST configuration.
show spanning-tree mst	Displays information about the MST protocol | Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 125.	The instance command inserts an entry into the VLAN-to-instance map that associates a set of VLANs to an MST instance. In addition to defining the MST topology, the VLAN-to-instance map is one of three parameters, along with the MST name and revision number, that identifies the switch's MST region. The no instance command removes specified entries from the VLAN-to-instance map. If the command does not provide a VLAN list, all entries are removed for the specified instance. The no instance and default instance commands function identically. Platform all Command Mode MST-Configuration Command Syntax instance mst_inst vlans v_range no instance mst_inst [vlans v_range] no default instance mst_inst [vlans v_range] Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 978.

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 148.	**spanning-tree bpduguard**
		
		The spanning-tree bpduguard command controls BPDU guard on the configuration mode interface. A BPDU guard-enabled port is disabled when it receives a BPDU packet. Disabled ports differ from blocked ports in that they are re-enabled only through manual intervention.
		
		The BPDU guard default setting for portfast ports is configured by the spanning-tree portfast bpduguard default command; BPDU guard is disabled by default on all non-portfast ports.
		
		• spanning-tree bpduguard enable enables BPDU guard on the interface.
		• spanning-tree bpduguard disable disables BPDU guard on the interface.
		
		The no spanning-tree bpduguard and default spanning-tree bpduguard commands restore the global BPDU guard setting on the configuration mode interface by removing the corresponding spanning-tree bpduguard command from *running-config*.
		
		Platform all
		Command Mode Interface-Ethernet Configuration
		                Interface-Port-Channel Configuration
		
		Command Syntax
		`spanning-tree bpduguard GUARD_ACTION`
		`no spanning-tree bpduguard`
		`default spanning-tree bpduguard`
		
		Parameters
		• *GUARD_ACTION* BPDU guard setting. Options include:
		
		— enabled BPDU guard is enabled on the interface.
		— disabled BPDU guard is disabled on the interface.
		
		Examples
		• These commands enable BPDU guard on Ethernet interface 5.
		
		`switch(config)#interface ethernet 5`
		`switch(config-if-Et5)#spanning-tree bpduguard enabled`
		`switch(config-if-Et5)#`
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 997.

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	⚠ **Caution** When disabling spanning tree on a VLAN using the no spanning-tree vlan *vlan-id* command, ensure that all switches and bridges in the VLAN have spanning tree disabled. You cannot disable spanning tree on some switches and bridges in a VLAN and leave it enabled on other switches and bridges in the same VLAN because switches and bridges with spanning tree enabled have incomplete information about the physical topology of the network. ⚠ **Caution** We do not recommend disabling spanning tree even in a topology that is free of physical loops. Spanning tree is a safeguard against misconfigurations and cabling errors. Do not disable spanning tree in a VLAN without ensuring that there are no physical loops present in the VLAN. Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 159.	**Warning** Disabling spanning tree is not recommended, even in topologies free of physical loops. Spanning tree guards against configuration mistakes and cabling errors. When disabling VLAN, ensure that there are no physical loops in the VLAN. **Important** When disabling spanning tree on a VLAN, ensure that all switches and bridges in the network disable spanning tree for the same VLAN. Disabling spanning tree on a subset of switches and bridges in a VLAN may have unexpected results because switches and bridges running spanning tree will have incomplete information regarding the network's physical topology. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1023.
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**switchport private-vlan trunk native vlan** To set the native VLAN for private VLAN promiscuous and isolated trunk ports, use the switchport private-vlan trunk native vlan command. To return to the default value, use the no form of this command. switchport private-vlan trunk native vlan *vlan-id* no switchport private-vlan trunk native vlan *vlan id* Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 177.	**switchport trunk native vlan** The switchport trunk native vlan command specifies the trunk mode native VLAN for the configuration mode interface. Interfaces in trunk mode associate untagged frames with the native VLAN. Trunk mode interfaces can also be configured to drop untagged frames. The default native VLAN for all interfaces is VLAN 1. The no switchport trunk native vlan and default switchport trunk native vlan commands restore VLAN 1 as the trunk mode native VLAN to the configuration mode interface by removing the corresponding switchport trunk native vlan command from *running-config*. Platform all Command Mode Interface-Ethernet Configuration Interface-Port-channel Configuration Command Syntax switchport trunk native vlan VLAN_ID no switchport trunk native vlan default switchport trunk native vlan Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 800.

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	 Cisco Nexus 7000 Series NX-OS Interfaces Command Reference (2013), at 179.	 Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 751.
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	 Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 42.	 Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1521.

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	This example shows how to clear all OSPF neighbor details for all OSPF instances:
	
	`switch# clear ip ospf neighbor *`
	
	This example shows how to clear all OSPF neighbor details for all neighbors on Ethernet interface 1/2 for OSPF instance 202:
	
	`switch# clear ip ospf 202 neighbor ethernet 1/2`
	
	Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 112.
		Examples
		• This command resets all OSPF neighbor statistics.
		
		`switch#clear ip ospf neighbor *`
		`switch#`
		
		• This command resets the OSPF neighbor statistics for the specified Ethernet 3 interface.
		
		`switch#clear ip ospf neighbor ethernet 3`
		`switch##`
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1420.

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	default-information originate (OSPF)
	
	To generate a default external route into an Open Shortest Path First (OSPF) routing domain, use the **default-information originate** command. To disable this feature, use the **no** form of this command.
	
	default-information originate [always] [route-map *map-name*]
	
	no default-information originate [always] [route-map *map-name*]
	
	Syntax Description
	always — (Optional) Specifies to always advertise the default route regardless of whether the route table has a default route.
	route-map *map-name* — (Optional) Specifies to advertise the default route if the route map is satisfied. The *map-name* argument can be any alphanumeric string up to 63 characters.
	
	Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 42.
		Examples
		• These commands will always advertise the OSPFv2 default route regardless of whether the switch has a default route configured.
		
		`switch(config)#router ospf 1`
		`switch((config-router-ospf)#default-information originate always`
		`switch(config-router-ospf)#show active`
		`router ospf 1`
		`    default-information originate always`
		
		• These commands advertise a default route with a metric of 100 and an external metric type of 1 if a default route is configured.
		
		`switch(config)#router ospf 1`
		`switch((config-router-ospf)#default-information originate metric 100 metric-type 1`
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1423.

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2

Effective date of registration: 11/13/2014 | **default-information originate (OSPFv3)**

To generate a default external route into an Open Shortest Path First version 3 (OSPFv3) routing domain, use the default-information originate command. To disable this feature, use the no form of this command.

default-information originate [always] [route-map map name]

no default-information originate [always] [route-map map-name]

Syntax Description:
always — (Optional) Specifies to always advertise the default route regardless of whether the route table has a default route.

route-map map-name — (Optional) Specifies to advertise the default route if the route map is satisfied. The map-name argument can be any alphanumeric string up to 63 characters.

Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 44. | Examples

• These commands will always advertise the OSPFv3 default route regardless of whether the switch has a default route configured.

```
switch(config)#ipv6 router ospf 1
switch(config-router-ospf3)#default-information originate always
switch(config-router-ospf3)#show active
ipv6 router ospf 1
    default-information originate always
```

• These commands configures OSPF area 1 as metric of 100 for the default route with an external metric type of Type 1.

```
switch(config)#ipv6 router ospf 1
switch(config-router-ospf3)#default-information originate metric 100 metric-type 1
switch(config-router-ospf3)#show active
ipv6 router ospf 1
    default-information originate metric 100 metric-type 1
switch(config-router-ospf3)#
```

Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1506. |

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**distance (EIGRP)** To allow the use of two administrative distances—internal and external—for the Enhanced Interior Gateway Routing Protocol (EIGRP) that could provide a better route to a node, use the **distance** command. To reset to default, use the **no** form of this command. distance *internal-distance external-distance* no distance **Syntax Description** *internal-distance* — Administrative distance for EIGRP internal routes. Internal routes are routes that are learned from another entity within the same autonomous system (AS). The distance can be a value from 1 to 255. The default value is 90. *external-distance* — Administrative distance for EIGRP external routes. External routes are routes for which the best path is learned from a source external to this autonomous system. The distance can be a value from 1 to 255. The default value is 170. **Defaults** *internal-distance*: 90 *external-distance*: 170 Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 61.	**distance bgp** The distance bgp command assigns an administrative distance to routes that the switch learns through BGP. Routers use administrative distances to select a route when two protocols provide routing information to the same destination. Distance values range from 1 to 255; lower distance values correspond to higher reliability. BGP routing tables do not include routes with a distance of 255. The distance command assigns distance values to external, internal, and local BGP routes: • external: External routes are routes for which the best path is learned from a neighbor external to the autonomous system. Default distance is 200. • internal: Internal routes are routes learned from a BGP entity within the same autonomous system. Default distance is 200. • local: Local routes are networks listed with a network router configuration command for that router or for networks that are redistributed from another process. Default distance is 200. The no distance bgp and default distance bgp commands restore the default administrative distances by removing the distance bgp command from *running-config*. Platform — all Command Mode — Router-BGP Configuration **Command Syntax** distance bgp external_dist [INTERNAL_LOCAL] no distance bgp default distance bgp **Parameters** • *external_dist* — distance assigned to external routes. Values range from 1 to 255. • INTERNAL_LOCAL — distance assigned to internal and local routes. Values for both routes range from 1 to 255. Options include: — <no parameter> *external_dist* value is assigned to internal and local routes. — *internal_dist local_dist* values assigned to internal (*internal_dist*) and local (*local_dist*) routes. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1583.

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	When you configure the **ip** command on an interface, the handling of proxy Address Resolution Protocol (ARP) requests changes (unless proxy ARP was disabled). Hosts send ARP requests to map an IP address to a MAC address. The GLBP gateway intercepts the ARP requests and replies to the ARP requests on behalf of the connected nodes. If a forwarder in the GLBP group is active, proxy ARP requests are answered using the MAC address of the first active forwarder in the group. If no forwarder is active, proxy ARP responses are suppressed. Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 256.	**Displaying ARP Entries** The show ip arp command displays ARP cache entries that map an IP address to a corresponding MAC address. The table displays addresses by their host names when the command includes the *resolve* argument. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1225.
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Expanded Community Lists** Expanded community lists are used to filter communities using a regular expression. Regular expressions are used to configure patterns to match community attributes. The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it will match the earliest part first. Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 274.	The order for matching using the * or + character is longest construct first. Nested constructs are matched from the outside in. Concatenated constructs are matched beginning at the left side. If a regular expression can match two different parts of an input string, it matches the earliest part first. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 107.
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	Static routes have a default administrative distance of 1. If you want a dynamic routing protocol to take precidence over a static route, you must configure the static route preference argument to be greater than the administrative distance of the dynamic routingprotocol. For example, routes derived with Enhanced Interior Gateway Routing Protocol (EIGRP) have a default administrative distance of 100. To have a static route that would be overridden by an EIGRP dynamic route, specify an administrative distance greater than 100. Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 337.	Static routes have a default administrative distance of 1. Assigning a higher administrative distance to a static route configures it to be overridden by dynamic routing data. For example, a static route with a distance value of 200 is overridden by OSPF intra-area routes with a default distance of 110. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1226.

Copyright Registration Information	Cisco	Arista
| Cisco NX-OS 6.2<br><br>Effective date of registration: 11/13/2014 | **is-type**<br><br>To configure the routing level for an instance of the Intermediate System-to-Intermediate System (IS-IS) routing process, use the is-type command. To reset the default value, use the no form of this command.<br><br>is-type {level-1 | level-1-2 | level-2}<br><br>no is-type {level-1 | level-1-2 | level-2}<br><br>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 407. | **is-type**<br><br>The is-type command configures the routing level for an instance of the IS-IS routing instance.<br><br>Platform          all<br>Command Mode    Router-IS-IS Configuration<br><br>Command Syntax<br>is-type *LAYER_VALUE*<br><br>Parameters<br>• *LAYER_VALUE*    layer value. Options include:<br>— level-1    The switch operates as a Level-1 (intra-area) router.<br>— level-2    The switch operates as a Level-2 (inter-area) router.<br><br>Example<br>• These commands configure Level 2 routing on interface Ethernet 5.<br><br>```
switch(config)#router isis Osiris
switch(config-router-isis)#is-type level-2
switch(config-router-isis)#
```<br><br>Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1691. |

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	isis hello-multiplier
	
	To specify the number of Intermediate System-to-Intermediate System (IS-IS) hello packets a neighbor must miss before the router should declare the adjacency as down, use the isis hello-multiplier command. To restore the default value, use the no form of this command.
	
	isis hello-multiplier *multiplier* {level-1
		level-2}
		
		no isis hello-multiplier {level-1
			level-2}
			
			Syntax Description
			multiplier — Integer value. Range: 3 to 1000. Default: 3.
			level-1 — Configures the hello multiplier independently for Level 1 adjacencies.
			level-2 — Configures the hello multiplier independently for Level 2 adjacencies.
			
			Command Default The default settings are as follows:
			• *multiplier*: 3
			• Level 1 and Level 2
			
			Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 423.
				isis hello-multiplier
				
				The isis hello-multiplier command specifies the number of IS-IS hello packets a neighbor must miss before the device should declare the adjacency as down.
				
				Each hello packet contains a hold time. The hold time informs the receiving devices how long to wait without seeing another hello from the sending device before considering the sending device down. The isis hello-multiplier command is used to calculate the hold time announced in hello packets by multiplying this number with the configured isis hello-interval.
				
				The no isis hello-multiplier and default isis hello-multiplier commands restore the default hello interval of 3 on the configuration mode interface by removing the isis hello-multiplier command from *running-config*.
				
				Platform — all
				Command Mode — Interface-Ethernet Configuration
				Interface-Loopback Configuration
				Interface-Port-channel Configuration
				Interface-VLAN Configuration
				
				Command Syntax
				`isis hello-multiplier` *factor*
				`no isis hello-multiplier`
				`default isis hello-multiplier`
				
				Parameters
				• *factor* — hello multiplier. Values range from 3 to 100; default is 3.
				
				Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1685.

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	isis priority
	
	To configure the priority of designated routers, use the isis priority command in interface configuration mode. To reset the default priority, use the no form of this command.
	
	isis priority *number-value* [level-1
		level-2]
		
		no isis priority [level-1
			level-2]
			
			Syntax Description
			number-value — Priority of a router and is a number from 0 to 127. The default value is 64.
			level-1 — (Optional) Sets the priority for Level 1 independently.
			level-2 — (Optional) Sets the priority for Level 2 independently.
			
			Defaults
			Priority of 64
			Level 1 and Level 2
			
			Command Modes
			Interface configuration
			
			SupportedUserRoles
			network admin
			vdc-admin
			
			Command History
			Release — Modification
			4.0(1) — This command was introduced.
			
			Usage Guidelines
			Priorities can be configured for Level 1 and Level 2 independently. Specifying the level-1 or level-2 keyword resets priority only for Level 1 or Level 2 routing, respectively.
			
			The priority is used to determine which router on a LAN will be the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority will become the DIS.
			
			In Intermediate System-to-Intermediate System (IS-IS), there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a router with a higher priority comes on line, it will take over the role from the current DIS. In the case of equal priorities, the highest MAC address breaks the tie.
			
			This command requires the Enterprise Services license.
			
			Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 433.
				isis priority
				
				The isis priority command configures IS-IS router priority for the configuration mode interface.
				
				The priority is used to determine which device will be the Designated Intermediate System (DIS). The device with the highest priority will become the DIS.
				
				In IS-IS, there is no backup designated router. Setting the priority to 0 lowers the chance of this system becoming the DIS, but does not prevent it. If a device with a higher priority comes on line, it will take over the role from the current DIS.
				
				The no isis priority and default isis priority commands restore the default priority (64) on the configuration mode interface.
				
				Platform — all
				Command Mode — Interface-Ethernet Configuration
				Interface-Loopback Configuration
				Interface-Port-channel Configuration
				Interface-VLAN Configuration
				
				Command Syntax
				`isis priority priority_level`
				`no isis priority`
				`default isis priority`
				
				Parameters
				• *priority_level* — priority level. Value ranges from 0 to 127. Default value is 64.
				
				Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1690.

EXHIBIT A TO CISCO'S OBJECTIONS AND RESPONSES TO ARISTA'S FIRST SET OF INTERROGATORIES

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**log-adjacency-changes (IS-IS)** To enable the router to send a syslog message when an Intermediate System-to-Intermediate System Intradomain Routing Protocol (IS-IS) neighbor goes up or down, use the log-adjacency-changes configuration mode command. To disable this function, use the no form of this command. log-adjacency-changes no log-adjacency-changes **Syntax Description** This command has no arguments or keywords. **Defaults** This command is enabled by default. **Command Modes** Router configuration VRF configuration **SupportedUserRoles** network admin vdc-admin **Command History** Release Modification 4.0(1) This command was introduced. **Usage Guidelines** The log-adjacency-changes command is on by default but only up/down (full/down) events are reported. **Examples** This example configures the router to send a syslog message when an IS-IS neighbor state changes: switch(config)# router isis switch(config-router)# log-adjacency-changes **Related Commands** Command Description feature isis Enables IS-IS on the router. router isis Enables IS-IS. Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 138.	**log-adjacency-changes (IS-IS)** The log-adjacency-changes command configures the switch to send syslog messages either when it detects IS-IS link state changes or when it detects that a neighbor has gone up or down. Log message sending is disabled by default. The default option is active when *running-config* does not contain any form of the command. Entering the command in any form replaces the previous command state in *running-config*. Platform all Command Mode Router-IS-IS Configuration Command Syntax log-adjacency-changes no log-adjacency-changes default log-adjacency-changes Examples • These commands configure the switch to send a syslog message when a neighbor goes up or down. switch(config)#router isis Osiris switch(config-router-isis)#log-adjacency-changes switch(config-router-isis)# • These commands configure not to log the peer changes. switch(config)#router isis Osiris switch(config-router-isis)#no log-adjacency-changes switch(config-router-isis)# Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1692.

EXHIBIT A TO CISCO'S OBJECTIONS AND RESPONSES TO ARISTA'S FIRST SET OF INTERROGATORIES

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**max-metric router-lsa (OSPF)** To configure the Open Shortest Path First (OSPF) protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their shortest path first (SPF) calculations, use the max-metric router-lsa command. To disable the advertisement of a maximum metric, use the no form of this command. max-metric router-lsa [external-lsa [*max-metric-value*]] [include-stub] [on-startup [*seconds* | wait-for bgp *tag*]] [summary-lsa [max-metric-value]] no max-metric router-lsa [external-lsa [*max-metric-value*]] [include-stub] [on-startup [*seconds* | wait-for bgp *tag*]] [summary-lsa [max-metric-value]]] external-lsa — Specifies the external LSA's. *max-metric-value* — (Optional) Specifies the max-metric values for external LSA's. The range is 1-65535. inlcude-stub — Advertises the max-metric for stub links. on-startup — (Optional) Configures the router to advertise a maximum metric at startup. *seconds* — (Optional) Maximum metric (in seconds) that is advertised for the specified time interval. The configurable range is from 5 to 86400 seconds. The default is 600 seconds. wait-for bgp *tag* — (Optional) Advertises a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds. summary-lsa — Specifies the summary LSA's. *max-metric-value* — (Optional) Specifies the max-metric value for summary LSAs. The range is from 1-65535. **Defaults** — Originates router link-state advertisements (LSAs) with normal link metrics. **Command Modes** — Router configuration / Router VRF configuration **SupportedUserRoles** — network-admin / vdc-admin **Command History** Release — Modification 4.0(1) — This command was introduced. Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 194.	**max-metric router-lsa (OSPFv2)** The max-metric router-lsa command allows the OSPF protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their SPF calculations. The no max-metric router-lsa and default max-metric router-lsa commands disable the advertisement of a maximum metric. Platform — all Command Mode — Router-OSPF Configuration Command Syntax max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY] no max-metric router-lsa [EXTERNAL] [STUB][STARTUP] [SUMMARY] default max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY] All parameters can be placed in any order. Parameters • EXTERNAL — advertised metric value. Values include: — <no parameter> Metric is set to the default value of 1. — external-lsa Configures the router to override the External LSA / NSSA-External metric with the maximum metric value. — external-lsa <1 to 16777215> The configurable range is from 1 to 0xFFFFFF. The default value is 0xFF0000. This range can be used with external LSA, summary LSA extensions to indicate the respective metric you want with the LSA. • STUB — advertised metric type. Values include: — <no parameter> Metric type is set to the default value of 2. — include-stub Advertises stub links in router-LSA with the max-metric value (0xFFFF). • STARTUP — limit scope of LSAs. Values include: — <no parameter> LSA can be translated — on-startup Configures the router to advertise a maximum metric at startup (only valid in no and default command formats). — on-startup wait-for-bgp Configures the router to advertise a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds. — on-startup <5 to 86400> Sets the maximum metric temporarily after a reboot to originate router-LSAs with the max-metric value. wait-for-bgp or an on-start time value is not included in no and default commands. • SUMMARY — advertised metric value. Values include: — <no parameter> Metric is set to the default value of 1. — summary-lsa Configures the router to override the summary LSA metric with the maximum metric value for both type 3 and type 4 Summary LSAs. — summary-lsa <1 to 16777215> Metric is set to the specified value. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1439.

Copyright Registration Information	Cisco	Arista
| Cisco NX-OS 6.2<br><br>Effective date of registration: 11/13/2014 | ```
BGP table version is 10, local router ID is 3.3.3.3
Status: s-suppressed, x-deleted, S-stale, d-dampened, h-history, *-valid, >-best
Path type: i-internal, e-external, c-confed, 1-local, a-aggregate, r-redist
Origin codes: i - IGP, e - EGP, ? - incomplete, | - multipath

   Network          Next Hop          Metric     LocPrf     Weight Path
* i200.0.1.100/32   201.0.25.1                     100        100 6553601 i
*>e                 201.0.13.1                                  0 6553601 i
* i200.0.2.100/32   201.0.25.1                     100        100 6553601 i
*>e                 201.0.13.1                                  0 6553601 i
*>1200.0.3.100/32   0.0.0.0                        100      32768 i
```<br><br>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 401. | ```
switch>show ip bgp neighbors 10.14.4.4 advertised-routes regexp _64502_
BGP routing table information for VRF default
Router identifier 172.24.78.191, local AS number 64498
Route status codes: s - suppressed, * - valid, > - active, E - ECMP head, e - ECMP
                                 S - Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
AS Path Attributes: Or-ID - Originator ID, C-LST - Cluster List, LL Nexthop - Link
Local Nexthop

     Network          Next Hop       Metric    LocPref Weight Path
* >  10.99.31.0/24    10.88.202.1    333       100      -     (64502 64503) 99 i
* >  10.99.41.0/24    10.88.202.1    333       100      -     (64502 64503) 99 i
* >  10.99.99.0/24    10.88.202.1    333       100      -     (64502 64504) 99 i
```<br><br>Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1637. |
| Cisco NX-OS 6.2

Effective date of registration: 11/13/2014 | **show ip bgp neighbors**

To display Border Gateway Protocol (BGP) neighbors, use the show ip bgp neighbors command

show ip bgp neighbors [addr advertised-routes | flap-statistics | paths | received-routes | routes [advertised | dampened | received]] | prefix] [vrf { all | vrf-name}]

Syntax Description
addr — IPv4 address. The format is x.x.x.x
advertised-routes — (Optional) Displays all the routes advertised to this neighbor.
flap-statistics — (Optional) Displays flap statistics for the routes received from this neighbor.
paths — (Optional) Displays AS paths learned from this neighbor.
received-routes — (Optional) Displays all the routes received from this neighbor.
routes — (Optional) Displays the routes received or advertised to or from this neighbor.
advertised — (Optional) Displays all the routes advertised for this neighbor.
dampened — (Optional) Displays all dampened routes received from this neighbor.
received — (Optional) Displays all the routes received from this neighbor.
prefix — (Optional) IPv6 prefix. The format is x.x.x.x/length.
vrf vrf-name — (Optional) Specifies the virtual router context (VRF) name. The name can be any case-sensitive, alphanumeric string up to 63 characters.
all — (Optional) Specifies all VRF.

Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 466. | **show ip bgp neighbors**

The show ip bgp neighbors command displays Border Gateway Protocol (BGP) and TCP session data for a specified IPv4 BGP neighbor, or for all IPv4 BGP neighbors if an address is not included.

Platform   all
Command Mode   EXEC

Command Syntax
show ip bgp neighbors [NEIGHBOR_ADDR] [VRF_INSTANCE]

Parameters
• NEIGHBOR_ADDR   location of neighbors. Options include:
— <no parameter>   command displays information for all IPv4 BGP neighbors.
— ipv4 addr   command displays information for specified neighbor.
• VRF_INSTANCE   specifies VRF instances.
— <no parameter>   displays routing table for context-active VRF.
— vrf vrf_name   displays routing table for the specified VRF.
— vrf all   displays routing table for all VRFs.
— vrf default   displays routing table for default VRF.

Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1632. |

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	Use the **ip ospf database** command to display information about different OSPF LSAs.
	
	When the link state advertisement is describing a network, the *link-state-id* argument can take one of two forms:
	
	• The network's IP address (such as Type 3 summary link advertisements and autonomous system external link advertisements).
	• A derived address obtained from the link state ID. (Note that masking a network links advertisement's link state ID with the network's subnet mask yields the network's IP address.)
	• When the link state advertisement is describing a router, the link state ID is always the described router's OSPF router ID.
	• When an autonomous system external advertisement (LS Type = 5) is describing a default route, its link state ID is set to Default Destination (0.0.0.0).
	
	This command requires the Enterprise Services license.
	
	Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 520.
		• *linkstate_id* Network segment described by the LSA (dotted decimal notation).
		  Value depends on the LSA type.
		  — When the LSA describes a network, the *linkstate-id* argument is one of the following:
		     The network IP address, as in Type 3 summary link advertisements and in autonomous system external link advertisements.
		     A derived address obtained from the link state ID. Masking a network links the advertisement link state ID with the network subnet mask yielding the network IP address.
		  — When the LSA describes a router, the link state ID is the OSPFv2 router ID of the router.
		  — When an autonomous system external advertisement (Type 5) describes a default route, its link state ID is set to the default destination (0.0.0.0).
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1454.

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	timers basic
	
	To adjust the Routing Information Protocol (RIP) network timers, use the timers basic command. To restore the default timers, use the no form of this command.
	
	timers basic *update invalid holddown flash*
	
	no timers basic
	
	Syntax Description — *update* — Rate (in seconds) at which updates are sent. The default is 30 seconds.
	invalid — Interval of time (in seconds) after which a route is declared invalid; it should be at least three times the value of the *update* argument. A route becomes invalid when no updates refresh the route. The route then enters into a *holddown* state where it is marked as inaccessible and advertised as unreachable. However, the route is still used to forward packets. The default is 180 seconds.
	holddown — Interval (in seconds) during which routing information regarding better paths is suppressed; it should be at least three times the value of the *update* argument. A route enters into a *holddown* state when an update packet is received that indicates that the route is unreachable. The route is marked as inaccessible and advertised as unreachable. However, the route is still used to forward packets. When holddown expires, routes advertised by other sources are accepted and the route is no longer inaccessible. The default is 180 seconds.
	flush — Amount of time (in seconds) that must pass before the route is removed from the routing table; the interval specified should be greater than the sum of the *invalid* argument plus the *holddown* argument. If it is less than this sum, the proper *holddown* interval cannot elapse, which results in a new route being accepted before the *holddown* interval expires. The default is 240 seconds.
	
	Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 1014.
		timers basic (RIP)
		
		The timers basic command configures the update interval, the expiration time, and the deletion time for routes received and sent through RIP. The command requires value declaration of all values.
		
		• The update time is the interval between unsolicited route responses. The default is 30 seconds.
		• The expiration time is initialized when a route is established and any time an update is received for the route. If the specified period elapses from the last time the route update was received, then the route is marked as inaccessible and advertised as unreachable. However, the route forwards packets until the deletion time expires. The default value is 180 seconds.
		• The deletion time is initialized when the expiration time has elapsed. On initialization of the deletion time, the route is no longer valid; however, it is retained in the routing table for a short time so that neighbors can be notified that the route has been dropped. Upon expiration of the deletion time, the route is removed from the routing table. The default is 120 seconds.
		
		The no timers basic and default timers basic commands return the timer values to their default values by removing the timers-basic command from *running-config*.
		
		Platform all
		Command Mode Router-RIP Configuration
		
		Command Syntax
		`timers basic` update_time expire_time deletion_time
		`no timers basic`
		`default timers basic`
		
		Parameters
		• *update time* — rate at which updates are sent.
		• *expire_time* — period a route is valid after it is established or updated. Must be greater than *update_time*.
		• *deletion_time* — interval after expiration when route is removed from routing table.
		  Parameter values are in seconds and range from 5 to 2,147,483,647.
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1671.

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**timers lsa-arrival (OSPF)** To set the minimum interval in which the software accepts the same link state advertisement (LSA) from Open Shortest Path First (OSPF) neighbors, use the timers lsa-arrival command. To return to the default, use the no form of this command. timers lsa-arrival *milliseconds* no timers lsa-arrival **Syntax Description** *milliseconds* — Minimum delay (in milliseconds) that must pass between acceptance of the same LSA arriving from neighbors. The range is from 10 to 600,000 milliseconds. The default is 1000 milliseconds. **Defaults** 1000 milliseconds **Command Modes** Router configuration VRF configuration **SupportedUserRoles** network-admin vdc-admin **Command History** Release Modification 4.0(1) This command was introduced. **Usage Guidelines** Use the timers lsa arrival command to configure the minimum interval for accepting the same LSA. The same LSA is an LSA instance that contains the same LSA ID number, LSA type, and advertising router ID. If an instance of the same LSA arrives sooner than the interval that is set, the software drops the LSA. We recommend that you keep the *milliseconds* value of the timers lsa-arrival command less than or equal to the neighbors' *hold interval* value of the timers throttle lsa command. This command requires the Enterprise Services license. **Examples** This example shows how to set the minimum interval for accepting the same LSA at 2000 milliseconds: switch(config)# router ospf 1 switch(config-router)# timers lsa-arrival 2000 Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 1016.	**timers lsa arrival (OSPFv2)** The timers lsa arrival command sets the minimum interval in which the switch accepts the same link-state advertisement (LSA) from OSPF) neighbors. The no timers lsa arrival and default timers lsa arrival commands restore the default maximum OSPFv2 path calculation interval to five seconds by removing the timers lsa arrival command from *running-config*. Platform all Command Mode Router-OSPF Configuration **Command Syntax** timers lsa arrival *lsa_time* no timers lsa arrival default timers lsa arrival **Parameters** • *lsa_time* OSPFv2 minimum interval (seconds). Values range from 1 to 600000 milliseconds. Default is 1000 milliseconds. **Example** • This command sets the minimum interval timer to ten milliseconds. switch(config)#router ospf 6 switch(config-router-ospf)#timers lsa arrival 10 switch(config-router-ospf)# Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1469.

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	Examples This example shows how to configure a router configured with the start, hold, and maximum interval values for the timers throttle spf command set at 5, 1000, and 90,000 milliseconds:
	`switch(config)# router ospf 1`
	`switch(config-router)# timers throttle spf 5 1000 90000`
	
	Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 1033-34.
		Example
		• This command displays a switch configured with the start, hold, and maximum interval values for the timers throttle spf command set at 5, 1,000, and 20,000 milliseconds, respectively.
		`switch(config)#router ospf 6`
		`switch(config-router-ospf)#timers spf 5 100 20000`
		`switch(config-router-ospf)#`
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1472.

Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	When using route reflectors, an AS is divided into clusters. A cluster consists of one or more route reflectors and a group of clients to which they re-advertise route information. Multiple route reflectors can be configured in the same cluster to increase redundancy and avoid a single point of failure. Each route reflector has a cluster ID. If the cluster has a single route reflector, the cluster ID is its router ID. If a cluster has multiple route reflectors, a 4-byte cluster ID is assigned to all route reflectors in the cluster. All of them must be configured with the same cluster ID so that they can recognize updates from other route reflectors in the same cluster. The bgp cluster-id command configures the cluster ID in a cluster with multiple route reflectors. Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference (2013), at 730.	cluster-id {cluster-id	cluster-ip-addr}—Configures the Route Reflector Cluster-ID (router, vrf). Range: 1 to 4294967295. You can enter the cluster identification as a 32-bit quantity or as an IP address. To remove the cluster ID, use the no form of this command. Together, a route reflector and its clients form a cluster. When a single route reflector is deployed in a cluster, the cluster is identified by the router ID of the route reflector. The cluster-id command is used to assign a cluster ID to a route reflector when the cluster has one or more route reflectors. Multiple route reflectors are deployed in a cluster to increase redundancy and avoid a single point of failure. When multiple route reflectors are configured in a cluster, the same cluster ID is assigned to all route reflectors. This allows all route reflectors in the cluster to recognize updates from peers in the same cluster and reduces the number of updates that need to be stored in BGP routing tables. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1549.
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Local Proxy ARP** You can use local Proxy ARP to enable a device to respond to ARP requests for IP addresses within a subnet where normally no routing is required. When you enable local Proxy ARP, ARP responds to all ARP requests for IP addresses within the subnet and forwards all traffic between hosts in the subnet. Use this feature only on subnets where hosts are intentionally prevented from communicating directly by the configuration on the device to which they are connected. Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 2-5.	**ip local-proxy-arp** The ip local-proxy-arp command enables local proxy ARP (Address Resolution Protocol) on the configuration mode interface. Local proxy ARP programs the switch to respond to ARP requests for IP addresses within a subnet where routing is not normally required. A typical local proxy arp application is supporting isolated private VLANs that communicate with each other by routing packets. The no ip local-proxy-arp and default ip local-proxy-arp commands disable local proxy ARP on the configuration mode interface by removing the corresponding ip local-proxy-arp command from running-config. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1276.

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	ipv6 nd managed-config-flag — Sets the managed address configuration flag in IPv6 router advertisements. ipv6 nd mtu — Sets the maximum transmission unit (MTU) size of IPv6 packets sent on an interface. ipv6 nd ns-interval — Configures the interval between IPv6 neighbor solicitation retransmissions on an interface. ipv6 nd other-config-flag — Configures the other stateful configuration flag in IPv6 router advertisements. Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 3-24.	**Router Advertisement Flag Configuration** The ipv6 nd managed-config-flag command configures the switch to set the *managed address configuration* flag in IPv6 router advertisements transmitted from the configuration mode interface. This bit instructs receptive hosts to use stateful address autoconfiguration. The ipv6 nd other-config-flag command configures the switch to set the *other stateful configuration* flag in IPv6 router advertisements transmitted from the configuration mode interface. This flag indicates the availability of autoconfiguration information, other than addresses, and that hosts should use stateful Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1329.
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	ipv6 nd reachable-time — Configures the amount of time that a remote IPv6 node is considered reachable after some reachability confirmation event has occurred. Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 3-24.	**ipv6 nd reachable-time** The ipv6 nd reachable-time command specifies the time period that the switch includes in the reachable time field of Router Advertisements (RAs) sent from the configuration mode interface. The reachable time defines the period that a remote IPv6 node is considered reachable after a reachability confirmation event. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1359.

Copyright Registration Information	Cisco	Arista
	Step 3 `max-metric router-lsa` `[external`-lsa `[max-metric-value]]` `[stub`-prefix-lsa] `[on`-`startup` `[seconds`	`wait-for-bgp` tag]] `[inter-area-prefix-lsa` `[max-metric-sumlsa]]` Example: `switch(config-router)`# max-metric router-lsa on-startup wait-for-bgp Configures a device that is running the OSPFv3 protocol to advertise a maximum metric so that other devices do not prefer the device as an intermediate hop in their SPF calculations. **Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 7-42.**	**max-metric router-lsa (OSPFv3)** The max-metric router-lsa command allows the OSPFv3 protocol to advertise a maximum metric so that other routers do not prefer the router as an intermediate hop in their SPF calculations. The no max-metric router-lsa and default max-metric router-lsa commands disable the advertisement of a maximum metric. Platform all Command Mode Router-OSPF3 Configuration Command Syntax `max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY]` `no max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY]` `default max-metric router-lsa [EXTERNAL] [STUB] [STARTUP] [SUMMARY]` All parameters can be placed in any order. Parameters • EXTERNAL advertised metric value. Values include: — <no parameter> Metric is set to the default value of 1. — external-lsa Configures the router to override the External LSA / NSSA-External metric with the maximum metric value. — external-lsa <1 to 16777215> The configurable range is from 1 to 0xFFFFFF. The default value is 0xFF0000. This range can be used with external LSA, summary LSA extensions to indicate the respective metric you want with the LSA. • STUB advertised metric type. Values include: — <no parameter> Metric type is set to the default value of 2. — include-stub Advertises stub links in router LSA with the max metric value (0xFFFF). • STARTUP limit scope of LSAs. Values include: — <no parameter> LSA can be translated — on-startup Configures the router to advertise a maximum metric at startup (only valid in no and default command formats). — on-startup wait-for-bgp Configures the router to advertise a maximum metric until Border Gateway Protocol (BGP) routing tables have converged or the default timer has expired. The default timer is 600 seconds. — on-startup <5 to 86400> Sets the maximum metric temporarily after a reboot to originate router-LSAs with the max-metric value. wait-for-bgp or an on-start time value is not included in no and default commands. • SUMMARY advertised metric value. Values include: — <no parameter> Metric is set to the default value of 1. — summary-lsa Configures the router to override the summary LSA metric with the maximum metric value for both type 3 and type 4 Summary LSAs. — summary-lsa <1 to 16777215> Metric is set to the specified value. Example • This command shows how to configure OSPFv3 to originate router LSAs with the maximum metric until BGP indicates that it has converged: `switch(config-router-ospf3)`#max-metric router-lsa on-startup wait-for-bgp `switch(config-router-ospf3)`# **Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1519.**
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014		

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2

Effective date of registration: 11/13/2014 | **IS-IS Overview**

IS-IS sends a hello packet out every configured interface to discover IS-IS neighbor routers. The hello packet contains information, such as the authentication, area, and supported protocols, which the receiving interface uses to determine compatibility with the originating interface. The hello packets are also padded to ensure that IS-IS establishes adjacencies only with interfaces that have matching maximum transmission unit (MTU) settings. Compatible interfaces form adjacencies, which update routing information in the link-state database through link-state update messages (LSPs). By default, the router sends a periodic LSP refresh every 10 minutes and the LSPs remain in the link-state database for 20 minutes (the LSP lifetime). If the router does not receive an LSP refresh before the end of the LSP lifetime, the router deletes the LSP from the database.

The LSP interval must be less than the LSP lifetime or the LSPs time out before they are refreshed.

IS-IS sends periodic hello packets to adjacent routers. If you configure transient mode for hello packets, these hello packets do not include the excess padding used before IS-IS establishes adjacencies. If the MTU value on adjacent routers changes, IS-IS can detect this change and send padded hello packets for a period of time. IS-IS uses this feature to detect mismatched MTU values on adjacent routers. For more information, see the "Configuring the Transient Mode for Hello Padding" section on page 9-21.

IS-IS Areas

You can design IS-IS networks as a single area that includes all routers in the network or as multiple areas that connect into a backbone or Level 2 area. Routers in a nonbackbone area are Level 1 routers that establish adjacencies within a local area (intra-area routing). Level 2 area routers establish adjacencies to other Level 2 routers and perform routing between Level 1 areas (inter-area routing). A router can have both Level 1 and Level 2 areas configured. These Level 1/Level 2 routers act as area border routers that route information from the local area to the Level 2 backbone area (see Figure 9-1).

Within a Level 1 area, routers know how to reach all other routers in that area. The Level 2 routers know how to reach other area border routers and other Level 2 routers. Level 1/Level 2 routers straddle the boundary between two areas, routing traffic to and from the Level 2 backbone area. Level1/Level2 routers use the attached (ATT) bit signal Level 1 routers to set a default route to this Level1/Level2 router to connect to the Level 2 area.

In some instances, such as when you have two or more Level1/Level 2 routers in an area, you may want to control which Level1/Level2 router that the Level 1 routers use as the default route to the Level 2 area. You can configure which Level1/Level2 router sets the attached bit. For more information, see the "Verifying the IS-IS Configuration" section on page 9-55.

Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 9-2. | **IS-IS Description**

IS-IS sends a hello packet out every configured interface to discover IS-IS neighbor routers. The hello packet contains information, which the receiving interface uses to determine compatibility with the originating interface. Compatible interfaces form adjacencies, which update routing information in the link-state database through link-state update messages (LSPs). If the router does not receive an LSP refresh before the end of the LSP lifetime, the device deletes the LSP from the database.

Terms of IS-IS Routing Protocol

The following terms are used when configuring IS-IS.

- NET and System ID – Each IS-IS instance has an associated network entity title (NET). The NET consists of the IS-IS system ID, which uniquely identifies the IS-IS instance in the area and the area ID.

- Designated Intermediate System – IS-IS uses a Designated Intermediate System (DIS) in broadcast networks to prevent each device from forming unnecessary links with every other device on the broadcast network. IS-IS devices send LSPs to the DIS, which manages all the link-state information for the broadcast network. You can configure the IS-IS priority that IS-IS uses to select the DIS in an area.

- IS-IS Areas – You can design IS-IS networks as a single area that includes all routers in the network or as multiple areas that connect into a backbone or Level 2 area. Routers in a nonbackbone area are Level 1 routers that establish adjacencies within a local area (intra-area routing). Level 2 area routers establish adjacencies to other Level 2 routers and perform routing between Level 1 areas (inter-area routing). A router can have both Level 1 and Level 2 areas configured.

- IS-IS Instances – Arista supports only one instance of the IS-IS protocol that run on the same node.

- LSP – Link state packet (LSP) can switch link state information. LSPs fall into two types: Level 1 LSPs and Level 2 LSPs. Level 2 devices transmit Level 2 LSPs; Level-1 devices transmit Level 1 LSPs; Level 1-2 devices transmit both Level 2 LSPs and Level 1 LSPs.

- Hello packets – Hello packets, can establish and maintain neighbor relationships.

- Overload Bit – IS-IS uses the overload bit to tell other devices not to use the local router to forward traffic but to continue routing traffic destined for that local router. Possible conditions for setting the overload bit the device is in a critical condition.

Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1674. |

EXHIBIT A TO CISCO'S OBJECTIONS AND RESPONSES TO ARISTA'S FIRST SET OF INTERROGATORIES

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**NET and System ID** Each IS-IS instance has an associated network entity title (NET). The NET is comprised of the IS-IS system ID, which uniquely identifies this IS IS instance in the area and the area ID. For example, if the NET is 47.0004.004d.0001.0001.0c11.1111.00, the system ID is 0000.0c11.1111.00 and the area is ID 47.0004.004d.0001. **Designated Intermediate System** IS IS uses a designated intermediate system (DIS) in broadcast networks to prevent each router from forming unnecessary links with every other router on the broadcast network. IS-IS routers send LSPs to the DIS, which manages all the link-state information for the broadcast network. You can configure the IS-IS priority that IS-IS uses to select the DIS in an area. Note No DIS is required on a point to point network. Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 9-3.	**Terms of IS-IS Routing Protocol** The following terms are used when configuring IS-IS. • NET and System ID – Each IS-IS instance has an associated network entity title (NET). The NET consists of the IS-IS system ID, which uniquely identifies the IS-IS instance in the area and the area ID. • Designated Intermediate System – IS-IS uses a Designated Intermediate System (DIS) in broadcast networks to prevent each device from forming unnecessary links with every other device on the broadcast network. IS-IS devices send LSPs to the DIS, which manages all the link-state information for the broadcast network. You can configure the IS-IS priority that IS-IS uses to select the DIS in an area. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1674.
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**Overload Bit** IS-IS uses the overload bit to tell other routers not to use the local router to forward traffic but to continue routing traffic destined for that local router. You may want to use the overload bit in these situations: • The router is in a critical condition. • Graceful introduction and removal of the router to/from the network. • Other (administrative or traffic engineering) reasons such as waiting for BGP convergence. Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 9-4.	• Overload Bit – IS-IS uses the overload bit to tell other devices not to use the local router to forward traffic but to continue routing traffic destined for that local router. Possible conditions for setting the overload bit the device is in a critical condition. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1674.

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	`isis hello-multiplier` *num* `[level-1 | level-2]`
	
	`Example:`
	`switch(config-if)# isis hello-multiplier 20`
	
	Specifies the number of IS-IS hello packets that a neighbor must miss before the router tears down an adjacency. The range is from 3 to 1000. The default is 3.
	
	Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 9-33.
		isis hello-multiplier
		
		The isis hello-multiplier command specifies the number of IS-IS hello packets a neighbor must miss before the device should declare the adjacency as down.
		
		Each hello packet contains a hold time. The hold time informs the receiving devices how long to wait without seeing another hello from the sending device before considering the sending device down. The isis hello-multiplier command is used to calculate the hold time announced in hello packets by multiplying this number with the configured isis hello-interval.
		
		The no isis hello-multiplier and default isis hello-multiplier commands restore the default hello interval of 3 on the configuration mode interface by removing the isis hello-multiplier command from *running-config*.
		
		Platform — all
		Command Mode — Interface-Ethernet Configuration / Interface-Loopback Configuration / Interface-Port-channel Configuration / Interface-VLAN Configuration
		
		Command Syntax
		`isis hello-multiplier` *factor*
		`no isis hello-multiplier`
		`default isis hello-multiplier`
		
		Parameters
		• *factor* — hello multiplier. Values range from 3 to 100; default is 3
		
		Examples
		• These commands configure a hello multiplier of 4 for VLAN 200.
		
		`switch(config)#interface vlan 200`
		`switch(config-if-V1200)#isis hello-multiplier 4`
		`switch(config-if-V1200)#`
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1685.

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	Step 9 `route-reflector-client`
	
	Example:
	`switch(config-router-neighbor-af)#`
	`route-reflector-client`
	
	Configures the device as a BGP route reflector and configures the neighbor as its client. This command triggers an automatic notification and session reset for the BGP neighbor sessions.
	
	Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 11-33.
		A route reflector is configured to re-advertise routes learned through IBGP to a group of BGP neighbors within the AS (its clients), eliminating the need for a fully meshed topology. The neighbor route-reflector-client command configures the switch to act as a route reflector and configures the specified neighbor as one of its clients. The bgp client-to-client reflection command enables client-to-client reflection.
		
		When using route reflectors, an AS is divided into clusters. A cluster consists of one or more route reflectors and a group of clients to which they re-advertise route information. Multiple route reflectors can be configured in the same cluster to increase redundancy and avoid a single point of failure. Each route reflector has a cluster ID. If the cluster has a single route reflector, the cluster ID is its router ID. If a cluster has multiple route reflectors, a 4-byte cluster ID is assigned to all route reflectors in the cluster. All of them must be configured with the same cluster ID so that they can recognize updates from other route reflectors in the same cluster. The bgp cluster-id command configures the cluster ID in a cluster with multiple route reflectors.
		
		Example
		• These commands configure the switch as a route reflector and the neighbor at 101.72.14.5 as one of its clients, and set the cluster ID to 172.22.30.101.
		`switch(config-router-bgp)#neighbor 101.72.14.5 route-reflector-client`
		`switch(config-router-bgp)#bgp cluster-id 172.22.30.101`
		`switch(config-router-bgp)#`
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1549.

Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	Static routes have a default administrative distance of 1. A router prefers a static route to a dynamic route because the router considers a route with a low number to be the shortest. If you want a dynamic route to override a static route, you can specify an administrative distance for the static route. For example, if you have two dynamic routes with an administrative distance of 120, you would specify an administrative distance that is greater than 120 for the static route if you want the dynamic route to override the static route. Cisco Nexus 7000 Series NX-OS Unicast Routing Configuration Guide, Release 6.x (2013), at 13-2.	Static routes have a default administrative distance of 1. Static routes with a higher administrative distance may be overridden by dynamic routing. For example, a static route with a distance of 200 is overridden by default OSPF intra-area routes (distance of 110). Route maps use tags to filter routes. Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1720.

Copyright Registration Information	Cisco	Arista
Cisco NX-OS 6.2 Effective date of registration: 11/13/2014	**clear ip igmp interface statistics** To clear the IGMP statistics for an interface, use the clear ip igmp interface statistics command. **clear ip igmp interface statistics** [*if-type if-number*] **Syntax Description** *if-type* — (Optional) Interface type. For more information, use the question mark (?) online help function. *if number* — (Optional) Interface or subinterface number. For more information about the numbering syntax for your networking device, use the question mark (?) online help function. **Defaults** — None **Command Modes** — Any command mode **SupportedUserRoles** — network admin network-operator vdc-admin vdc-operator **Command History** Release — Modification 4.0(3) — This command was introduced. **Usage Guidelines** — This command does not require a license. **Examples** This example shows how to clear IGMP statistics for an interface: switch# clear ip igmp interface statistics ethernet 2/1 switch# **Related Commands** Command — Description show ip igmp interface — Displays information about IGMP interfaces. Cisco Nexus 7000 Series NX-OS Multicast Routing Command Reference (2013), at 6.	**clear ip igmp statistics** The clear ip igmp statistics command resets IGMP transmission statistic counters for the specified interface. Platform — all Command Mode — Privileged EXEC **Command Syntax** clear ip igmp statistics [*INTF_ID*] **Parameters** • *INTF_ID* — interface name. Options include: — <no parameter> all interfaces. — interface ethernet *e_num* Ethernet interface specified by *e_num*. — interface loopback *l_num* Loopback interface specified by *l_num*. — interface management *m_num* Management interface specified by *m_num*. — interface port-channel *p_num* Port-channel interface specified by *p_num*. — interface vlan *v_num* VLAN interface specified by *v_num*. — interface xlan *vx_num* VXLAN interface specified by *vx_num*. **Examples** • This command resets IGMP transmission statistic counters on Ethernet 1 interface. switch#clear ip igmp statistics interface ethernet 1 switch# Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1794.

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	ip igmp snooping last-member-query-interval
	
	To configure a query interval in which the software removes a group, use the **ip igmp snooping last-member-query-interval** command. To reset the query interval to the default, use the **no** form of this command.
	
	ip igmp snooping last-member-query-interval [*interval*]
	no ip igmp snooping last-member-query-interval [*interval*]
	
	Syntax Description *interval* — Query interval in seconds. The range is from 1 to 25. The default is 1.
	
	Defaults The query interval is 1.
	
	Command Modes VLAN configuration (config-vlan) until Cisco NX-OS Release 5.1.
	Configure VLAN (config-vlan-config) since Cisco NS-OS Release 5.1(1). You cannot configure this command in the VLAN configuration mode in Cisco Release NX-OS 5.1 and higher.
	
	SupportedUserRoles network-admin / vdc-admin
	
	Command History
	Release: NX-OS 5.1(1) — Modification: The mode to configure this command on a VLAN changed to the configure VLAN mode (config-vlan-config)#. You can no longer configure this command in the VLAN configuration mode (config-vlan)#.
	Release: 4.0(1) — Modification: This command was introduced.
	
	Usage Guidelines This command does not require a license.
	See the Layer2 Command Reference Guide for information on entering the Configure VLAN mode by using the vlan configuration command.
	
	Examples This example shows how to configure a query interval in which the software removes a group:
	`switch(config)# vlan configuration 10`
	`switch(config-vlan-config)# ip igmp snooping last-member-query-interval 3`
	`switch(config-vlan-config)#`
	
	This example shows how to reset a query interval to the default:
	`switch(config)# vlan configuration 10`
	`switch(config-vlan-config)# no ip igmp snooping last-member-query-interval`
	`switch(config-vlan-config)#`
	
	Cisco Nexus 7000 Series NX-OS Multicast Routing Command Reference (2013), at 86.
		ip igmp last-member-query-interval
		
		The ip igmp last-member-query-interval command configures the switch's transmission interval for sending group-specific or group-source-specific query messages from the configuration mode interface.
		
		When a switch receives a message from a host that is leaving a group it sends query messages at intervals set by this command. The ip igmp startup-query-count specifies the number of messages that are sent before the switch stops forwarding packets to the host.
		
		If the switch does not receive a response after this period, it stops forwarding traffic to the host on behalf of the group, source, or channel.
		
		The no ip igmp last-member-query-interval and default ip igmp last-member-query-interval commands reset the query interval to the default value of one second by removing the ip igmp last-member-query-interval command from *running-config*.
		
		Platform: all
		Command Mode: Interface-Ethernet Configuration / Interface-Port-Channel Configuration / Interface-VLAN Configuration
		
		Command Syntax
		`ip igmp last-member-query-interval` *period*
		`no ip igmp last-member-query-interval`
		`default ip igmp last-member-query-interval`
		
		Parameters
		• *period* — transmission interval (deciseconds) between consecutive group-specific query messages.
		Value range: 10 (one second) to 317440 (8 hours, 49 minutes, 4 seconds). Default is 10 (one second).
		
		Example
		• This command configures the last member query interval of 6 seconds for VLAN interface 4.
		`switch(config)#interface vlan 4`
		`switch(config-if-Vl4)#ip igmp last-member-query-interval 60`
		`switch(config-if-Vl4)#`
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1799.

Copyright Registration Information	Cisco	Arista

Cisco NX-OS 6.2
 	
Effective date of registration: 11/13/2014
	ip igmp snooping startup-query-count
	
	To configure the number of queries sent at startup, use the **ip igmp snooping startup-query-count** command. To return to the default settings, use the **no** form of this command.
	
	ip igmp snooping startup-query-count *value*
	no ip igmp snooping startup-query-count *value*
	
	Syntax Description *value* Count value. The range is from 1 to 10.
	
	Defaults None
	
	Command Modes VLAN configuration (config-vlan)
	
	SupportedUserRoles network-admin
	vdc-admin
	
	Command History Release Modification
	NX-OS 5.1(1) This command was introduced.
	
	Usage Guidelines This command does not require a license.
	
	Examples This example shows how to configure the number of queries sent at startup:
	switch(config)# **vlan configuration 10**
	switch(config-vlan-config)# **ip igmp snooping startup-query-count** 4
	switch(config-vlan-config)#
	
	Related Commands Command Description
	show ip igmp snooping Displays IGMP snooping information.
	
	Cisco Nexus 7000 Series NX-OS Multicast Routing Command Reference (2013), at 104.
		ip igmp snooping querier startup-query-count
		
		The ip igmp snooping querier startup-query-count command configures the global *startup query count* value. The *startup query count* specifies the number of query messages that the querier sends on a VLAN during the *startup query interval* (ip igmp snooping querier startup-query-interval).
		
		When snooping is enabled, the group state is more quickly established by sending query messages at a higher frequency. The *startup-query-interval* and *startup-query-count* parameters define the startup period by defining the number of queries to be sent and transmission frequency for these messages.
		
		VLANs use the global *startup query count* value when they are not assigned a value (ip igmp snooping vlan querier startup-query-count). VLAN commands take precedence over the global value. The default global value is specified by the robustness variable (ip igmp snooping robustness-variable).
		
		The no ip igmp snooping querier startup-query-count and default ip igmp snooping querier startup-query-count commands restore the default *startup-query-count* value by removing the corresponding ip igmp snooping querier startup-query-count command from *running-config*.
		
		Platform all
		Command Mode Global Configuration
		
		Command Syntax
		ip igmp snooping querier startup-query-count *number*
		no ip igmp snooping querier startup-query-count
		default ip igmp snooping querier startup-query-count
		
		Parameters
		• *number* global startup query count. Value ranges from 1 to 3.
		
		Example
		• These commands configure the global startup query count value of 2, then displays the status of the snooping querier.
		switch(config)#ip igmp snooping querier startup-query-count 2
		switch(config)#show ip igmp snooping querier status
		  Global IGMP Querier status
		-- 	
		admin state : Disabled
		source IP address : 0.0.0.0
		query-interval (sec) : 125.0
		max-response-time (sec) : 10.0
		querier timeout (sec) : 255.0
		last-member-query-interval (sec) : 1.0
		last-member-query-count : 2 (robustness)
		startup-query-interval (sec) : 31.25 (query-interval/4)
		startup-query-count : 2
		
		Vlan Admin IP Query Response Querier Operational Ver
		     State Interval Time Timeout State
		-- 	
		1 Disabled 0.0.0.0 125.0 10.0 255.0 Non-Querier v2
		100 Disabled 0.0.0.0 125.0 10.0 255.0 Non-Querier v2
		101 Disabled 0.0.0.0 125.0 10.0 255.0 Non-Querier v2
		switch(config)#
		
		Arista User Manual v. 4.14.3F – Rev. 2 (October 2, 2014), at 1813.